

THE DEEP SPACE 1 EXTENDED MISSION: CHALLENGES IN PREPARING FOR AN ENCOUNTER WITH COMET BORRELLY

Marc D. Rayman
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109 USA

Following the successful conclusion of its 11-month primary mission in September 1999, Deep Space 1 embarked on an ambitious extended mission. The spacecraft is using its ion propulsion system to help deliver it to an encounter with comet 19P/Borrelly in September 2001. Along the way, the very small operations team has faced numerous challenges with an aging and debilitated spacecraft that was not designed for a comet encounter. The progress in reaching comet Borrelly is described as are the plans for attempting to return science data from this high-risk conclusion to the extended mission.

INTRODUCTION

Deep Space 1 (DS1), the first mission of NASA's New Millennium Program (NMP), was launched on October 24, 1998. The primary mission was devoted to the testing and evaluation of 12 technologies selected by NMP. The technology experiments composing DS1's payload were selected on the bases of their importance to future space and Earth science programs, the significant advancements they offer over current state-of-the-art, the high risk they present to the first user, and the need for in-flight testing to reduce that risk.

In addition to its objectives of testing advanced, high-risk technologies, DS1 was intended to explore the limits of schedule and cost for development. Launch was 39 months after the beginning of pre-phase A. At the end of the primary mission in September 1999, the total project cost was under \$150 M (in real-year dollars). This includes development, operations, and the launch service. It does not include the development cost of some of the technologies that composed the payload, but it does include the integration of all the technologies. Despite the very short schedule and small budget, DS1 met or exceeded all of the mission

success criteria. The results of the technology testing and the other activities of the primary mission have been described elsewhere.^{1,2}

In September 1999, DS1 began an extended mission. The two-year extension will be accomplished for less than \$10 M (with \$1.4 M of that for the DS1 Science Team). For the second year of the extended mission, 18 full-time equivalents conducted all the work, including spacecraft operations, new flight software development and loading, encounter planning and testing, navigation, telemetry control and data management, Deep Space Network (DSN) scheduling, project management, coffee club maintenance, etc.

In November 1999, the spacecraft's stellar reference unit (SRU) failed, depriving it of 3-axis attitude knowledge and control. This was a critical loss, and with the primary mission already successful and complete, termination of the extended mission was given serious consideration. Nevertheless, the project undertook a very ambitious two-phase, seven-month recovery using new software and new operational methods. Rayman and Varghese³ describe the details of the loss of the SRU, the complex and successful rescue that followed, and the progress of mission operations through late September 2000.

CRUISE OPERATIONS

Throughout the time the recovery was underway, sustained thrusting with the

Copyright © 2001 by the American Institute of Aeronautics and Astronautics, Inc. The U.S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for Governmental purposes. All other rights are reserved by the copyright owner.

spacecraft's ion propulsion system (IPS) was not possible. Thrusting resumed at the end of June 2000, targeting an encounter with comet 19P/Borrelly in September 2001. One component of the recovery from the failed SRU was the use of the visible charge coupled device (CCD) in the miniature integrated/camera spectrometer (MICAS) to track a preselected bright star. This forced the design of the trajectory to Borrelly to use piecewise inertially fixed thrust segments, each aligned with a star (known as a "thrustar") that satisfied criteria for serving as an attitude reference.

When the IPS is thrusting, the attitude control system (ACS) uses it to control two spacecraft axes by gimbaling the ion thruster in a mode referred to as thrust vector control (TVC). The roll around the thrust axis is controlled with the hydrazine-based reaction control system (RCS), and all three axes are controlled with the RCS when the IPS is not in use; the spacecraft does not have reaction wheels. The RCS is also used for spacecraft turns even when operating in TVC and, in certain cases described below, for trajectory correction maneuvers (TCMs).

Because of the long time the spacecraft was unable to operate in TVC and the additional expenditure of hydrazine during the initial testing and use of the new system to recover from the loss of the SRU, the hydrazine margin for the remainder of the mission was very low. At launch, the spacecraft carried 31.1 kg of hydrazine; at the end of September 2000, it had approximately 9 kg remaining.

Projections for hydrazine consumption were consistent with exhausting the supply before the September 2001 encounter, so several measures were implemented to reduce the expenditure of hydrazine.

The most significant modification that decreased the hydrazine consumption was the use of IPS thrusting whenever possible, in order to take advantage of the hydrazine savings in TVC mode. When IPS thrusting was most beneficial to the trajectory, thrusting was at the highest achievable throttle level (determined by the solar array power generation and the power needs for all systems except the IPS); otherwise, it was at a low throttle level. At this low level, referred to as "impulse power," the IPS thrust is

22.4 mN, thus accelerating the spacecraft by nearly 5 m/s/day. Targeting for the Borrelly encounter included all thrusting, whether at the optimal level or at impulse power.

Other methods to reduce hydrazine consumption are described later.

The first superior conjunction since DS1's launch occurred in November 2000 near the peak of solar cycle #23. Solar scintillation effects were expected to be significant for X-band, DS1's prime communications frequency, at Sun-Earth-probe (SEP) angles less than about 3° . Sequencing was designed so that no telecommunications within 5° of the solar limb would be needed.

The trajectory had been designed in such a way that IPS thrusting during conjunction was not required, because any problems that might have prevented IPS operation would have been difficult or impossible to correct during that period. Because of the need to continue thrusting to conserve hydrazine, the trajectory did assume operation at impulse power throughout conjunction. An attitude (and corresponding thrustar) was derived that helped the trajectory and would allow the unarticulated high-gain antenna (HGA) to be pointed in the vicinity of Earth throughout conjunction. This enabled opportunities to verify the spacecraft's health using the HGA or a co-boresighted low-gain antenna, should telecommunications be possible, and to study solar coronal effects on the radio frequency signals. The simultaneous use of X-band and K_a -band during two DSN passes with $SEP < 2^\circ$ enhanced these studies.⁴

The spacecraft turned to the conjunction thruster on October 30, 2000 and remained locked throughout the conjunction period, which ended on November 28. On November 7, 2000 DS1 reached its maximum geocentric range of 2.36 AU. The spacecraft was behind the solar disk from November 11 to November 13. During DSN coverage on November 14, with $SEP < 0.4^\circ$, Doppler data were noisy but adequate to confirm that the IPS had been thrusting most of the time since the beginning of the conjunction period. Detection of a sideband (described below) also confirmed that ACS was still locked to the thruster. During conjunction the IPS operated for 699.74 hours, its longest uninterrupted period of the primary or extended

mission.

From October 18, 2000 through January 2, 2001, the IPS operated at impulse power. Thrusting at the highest achievable throttle level resumed on January 2. Trajectory optimization studies had demonstrated that the use of impulse power at different times would have been more efficient in terms of xenon consumption. That metric was not the important one for DS1 however, as xenon was not a limiting resource. Instead, reducing the susceptibility to unplanned loss of IPS thrust was the dominant criterion in developing the thrust plan.⁵

In order to increase the probability of acquiring remote sensing data during the encounter with comet Borrelly, new flight software was designed, developed, tested, and integrated in late 2000 and early 2001. The 4 megabytes of new software were transmitted to the spacecraft from March 5 through March 8. The trajectory design had accounted for the time the spacecraft would spend thrusting in the attitude required to keep the HGA on Earth for the upload, installation, and verification of the new software as well as for the short time the IPS would have to be off. The spacecraft has only one central computer, so the reboot required to run the new software caused the spacecraft to enter one of its safe modes. This procedure had been executed successfully three times during the primary mission to increase the capability to test some of the technologies and once during the extended mission to recover from the loss of the SRU. The March 2001 procedure was more complicated, however, because of the need to conserve hydrazine during the safing recovery. Without a functional SRU, a safing causes the spacecraft to lose its full 3-axis attitude knowledge and control, the restoration of which requires significant ground intervention.

The computer reset to begin running the new software was commanded on March 13. Following the interactive coning procedure to bring the HGA to Earth-point,⁵ the IPS was restarted so TVC operation could resume. In the absence of an attitude reference star, the spacecraft relied on its inertial measurement units (IMUs) and the single Sun sensor assembly (SSA) to hold attitude, and the received signal strength at the DSN was used to sense spacecraft attitude drift. Short rotations about the Sun-spacecraft line were commanded

as necessary to maintain adequate HGA pointing. In normal operations without the SRU, ACS estimates the IMU bias (that is, the error in the IMU's rate measurement) for each axis using SSA and MICAS data. The reboot destroyed the bias estimates, however, and the recovery was complicated by the consequent significant attitude drift. If there had been adequate hydrazine, the system to search for an attitude reference star could have been invoked so that once a star were located, the spacecraft attitude would remain stable. This system had been used successfully during the recovery in June 2000, but the hydrazine consumption from the frequent small search turns was no longer affordable.

The spacecraft attitude could be estimated by combining the knowledge that the HGA was approximately pointed to Earth with telemetry of the Sun's position in spacecraft body coordinates, as measured by the SSA. Deep images returned from MICAS were analyzed to identify stars so that the spacecraft's attitude could be determined more accurately. The positive identification of stars in 3 images allowed the drift rate, and thus the IMU biases, to be estimated. With all these data, the on-board bias estimates were updated, and a turn to a known reference star was commanded, along with an estimated quaternion and the reenabling of the system to lock to a star. Following the first attempt of this, images revealed that the spacecraft had locked to a star 2° from the desired one. With the attitude stable, however, it was more straightforward to correct the spacecraft's quaternion.

The end of deterministic thrusting was reached on May 1, 2001, at which time the spacecraft was on a trajectory that would intercept comet Borrelly without further IPS thrusting. It remained essential to continue thrusting however to conserve hydrazine, so a thrust plan was adopted in which nearly all thrusting in any attitude was canceled by antiparallel thrusting at another time. Thus, every one or two weeks, the spacecraft would switch between impulse power thrusting toward the north ecliptic pole and thrusting toward the south ecliptic pole. Some of the time that the HGA was Earth-pointed the spacecraft thrustured approximately prograde and other times it would thrust retrograde. Distant encounter targeting was controlled by adjusting the throttle level in

some of these preplanned attitudes or changing the attitude slightly. In all cases, the spacecraft needed to be locked to an attitude reference star.

On September 1, 2001, DS1 had 580 days of operation on the IPS for thrusting to reach encounter targets, thrusting to allow TVC operation, and technology tests (for the IPS itself and some other technologies). It had consumed about 63.5 kg of xenon and imparted about 3.8 km/s to the spacecraft.

Prior to the need to reduce hydrazine consumption, one DSN pass each week was used to return engineering telemetry and conduct any necessary commanding. Because the RCS was used for turns and for locking the spacecraft to a new attitude reference star, a revised schedule of two DSN passes every three weeks was adopted. This allowed adequate contact with the spacecraft even if a scheduled DSN pass were lost and allowed the spacecraft to expend less hydrazine. As a small secondary benefit during deterministic thrusting, it permitted the spacecraft to achieve a higher thrusting duty cycle by not losing the time required to be in the Earth-point attitude every week.

Although the IPS was used to control two axes of the attitude, RCS still controlled the roll around the thrust axis. A further small reduction of hydrazine consumption was achieved by increasing the deadband around that axis from 1° to 2°. Still larger deadbands would have allowed occasional insolation on one of MICAS' radiators, thus compromising image quality; because MICAS was used for attitude control, this was unacceptable.

In most normal operating modes, when ACS was not locked to a reference star, a timer incremented. The timer was reset every time a picture of the reference star provided an update to the on-board attitude estimator. If the timer exceeded a preset value, a condition designated "celestial inertial reference loss" (CIRL), fault protection would trigger a safing. (When the SRU was operating, fault protection triggered the same response in the absence of good SRU data.)

The cost of a safing event and the subsequent recovery became significant when the hydrazine margin became small. In addition, before enough operational experience had been

accumulated with the system used to lock to a reference star, there was concern that it might lose lock frequently. Without data from MICAS, errors in the estimates of the IMU biases would cause the spacecraft to rotate around the Sun-spacecraft line. Drift around other axes would not occur, as ACS would continue to use SSA data to keep the Sun at the specified location in spacecraft body coordinates.

The system to lock to an attitude reference star included a mosaicking capability to find the star. The system was configured so that upon completing a turn or losing lock, it would be enabled to execute a 3×3 mosaic (each element of the mosaic corresponding to one MICAS 13-mrad field of view, with some overlap from one station to the next). When first operated, if it did not locate a star satisfying its criteria, the spacecraft would execute a short turn and commence another mosaic starting from this new attitude. Mosaic turns consumed a significant amount of hydrazine, however, so the system was altered to allow it to execute only one 3×3 mosaic. If that did not culminate in the location of an acceptable star, the spacecraft would switch to simply looking where it was pointing, allowing the effect of the IMU biases to change the attitude gradually until a satisfactory star moved into the field of view. By September 1, 2001 there had been only one instance of the spacecraft failing to lock to the targeted reference star following a turn, even for turns as large as 180°. That case is described below.

In September 2000 new software was dynamically linked to the running software. The short additional code monitored the CIRL timer, and when it reached a preset value, the software activated a stored command sequence. In this way, a less drastic response could be triggered before fault protection's CIRL-induced safing. The sequence that was activated for this "pre-CIRL" response was designed to avoid excessive expenditure of hydrazine and alert the operations team that there might be a problem.

If ACS lost its reference star and found a different one, for a variety of reasons the lock to the new star might not be solid. One situation that was considered important to avoid was the system losing lock, mosaicking, eventually finding a new star, and then losing lock again.

Each time it found a star, the single 3×3 mosaic would be reenabled for the subsequent attempt to lock. Thus, a series of weak locks on stars could lead to excessive hydrazine consumption. To prevent this, the sequence triggered by the pre-CIRL response disabled mosaicking following a subsequent loss of lock.

Once the spacecraft lost its reference star, sequenced turns would be unlikely to terminate within the 3×3 mosaic capture range of the planned reference star. Therefore, the pre-CIRL response sequence deleted sequences that would turn the spacecraft from Earth-point to a thruster. In this way, if the loss of lock occurred at the thrust attitude, the spacecraft would attempt to turn back to Earth (accepting the small chance that the 3×3 mosaic would find a star) but would not subsequently turn to a thruster. If the loss occurred during a DSN pass, the spacecraft would remain with its HGA pointed to Earth. In general, real-time monitoring during DSN passes was not feasible because of the very small size of the operations team. This automated prevention of a turn would allow the operations team to intervene before the spacecraft could turn back to an attitude in which communications would be difficult or impossible for up to two weeks.

Throughout the mission, when DS1 was at the thrust attitude, occasional short DSN passes were used to return low-rate telemetry, if possible, through a low-gain antenna (LGA) or to measure the Doppler shift. For DSN passes in which return of telemetry was not possible (usually because a 34-m station was used), the spacecraft would transmit a carrier with a subcarrier whose frequency depended upon whether the pre-CIRL response had been triggered. Normally the sideband would be at 35 kHz, but the pre-CIRL sequence would change it to 20 kHz. Detection of the 20 kHz signal would alert the operations team that the spacecraft might need assistance when it turned back to Earth-point.

The first time the pre-CIRL response was triggered was on July 15, 2001. As part of the strategy of controlling the trajectory, the spacecraft turned from one thruster near the north ecliptic pole to another and changed the IPS throttle level. It failed to lock to that star for reasons that were later determined to be related to scattered light in MICAS and the unusually

long imaging integration time required to lock to this dim reference star. After drifting 8° around the Sun-spacecraft line, the spacecraft locked to another star and remained stable. The 20 kHz tone was detected on July 18, allowing the operations team to prepare a response for the attempted turn to point the HGA to Earth on July 20.

The altered attitude of the spacecraft changed the projection of the IPS thrust vector onto the Earth-spacecraft line, so the Doppler data allowed the spacecraft attitude to be estimated. When the spacecraft turned to Earth, the received signal strength was used to estimate its pointing error, yielding a result consistent with the prediction based on the Doppler data. Working with images from MICAS, the operations team was able to correct the attitude and lock it to the planned star during the scheduled DSN pass. The extra hydrazine consumed as a result of the loss of lock, including the activities to relock it, was about 0.06 kg, or 1% of the remaining supply.

In August 2001 the spacecraft lost lock on its reference star again. In this third loss of lock since the June 2000 recovery, just as in the first (in July 2000), the loss occurred while it was tracking a star. In both cases the problem was triggered by an increased flux of solar protons flooding the MICAS CCD, thus confusing the star detection system. The recovery procedure was similar to the July 2001 recovery. While the recoveries were successful, the losses of lock illustrate the fragility of spacecraft operations without the star tracker.

ENCOUNTER PLANS

Closest approach to comet Borrelly will occur on September 22, 2001 at 1.36 AU from the Sun, 8 days after the comet's perihelion. The comet will be at a solar elongation of 63° , allowing simultaneous observations from Earth. In addition to ground-based facilities, the Hubble Space Telescope and Chandra X-ray Observatory are scheduled to observe the comet at or near the time of encounter. The spacecraft will be targeted to pass approximately 2000 km from the nucleus on the Sun-nucleus line between the Sun and nucleus with $v_\infty = 16.5$ km/s.

The encounter will be on the thirteenth recorded apparition of Borrelly. (Unfavorable

orbital conditions prevented the comet from being recovered on the fifth and sixth returns after its 1904 discovery.) With a period of 6.9 years, this Jupiter-family comet has been extensively studied by many investigators; it is moderately active with a well-defined coma and tails. A'Hearn *et al.*⁶ have identified Borrelly as a member of a compositional class of comets depleted in C-chain molecules but not in NH (all with respect to OH). Observations by Lamy *et al.*⁷ and by others suggest a nucleus with an equivalent spherical radius of about 2.5 km (assuming a geometric albedo of 0.04) and a rotation period of about 25 hours.

The principal measurements to be attempted at encounter are:

- ion and electron energy and angle spectra and ion mass/charge;
- magnetic field;
- panchromatic images of the nucleus and coma, with a target of a 50-pixel-diameter image of the nucleus; and
- infrared spectra of the nucleus.

The first measurements will be accomplished with the plasma experiment for planetary exploration (PEPE).⁸ PEPE will measure composition up to about 100 amu/e by directing ions into a cylinder with an electric field whose amplitude varies linearly with axial position. The resulting harmonic motion of an ion is independent of energy, so its travel time is a direct measure of mass/charge. During the primary mission, this time-of-flight (TOF) section was operated with the potential held at -8 kV on one end and +8 kV on the other. PEPE experienced an internal discharge in November 1999 (it is only fortuitous that it occurred in the same month as the SRU failure; the events are not causally related) that prevents it from holding the full positive potential. Testing on the spacecraft in November and December 2000 showed that the instrument could be operated with the TOF power supplies at -11 kV and +5 kV, and the data quality remained high.

PEPE can collect far more data than the data transfer system can accommodate, so its internal software bins the data in angle, energy, and mass/charge. Software to set this binning for

the encounter was uploaded in March 2001 and verified during both in-flight encounter simulations (described below). These tests provided further PEPE calibration by observing the solar wind.

As part of the primary mission's technology testing, a suite of sensors was included to measure the effects of the IPS on the spacecraft and space environment.⁹ These IPS diagnostic sensors (IDS) include two 3-axis fluxgate magnetometers which will be used for the cometary magnetic field measurements. IDS detected a magnetic field when the spacecraft encountered (9969) Braille.¹⁰ For the Braille encounter and IPS testing, magnetic field measurements could be made only intermittently. In June 2001 the IDS computer was reprogrammed to enable uninterrupted 20 Hz sampling of the magnetometers. To accommodate the data volume with the fixed internal IDS buffer size, the allocation for the plasma wave sensor data was reduced. Plasma wave data also will be recorded during the encounter.

The remote sensing data acquired with MICAS will rely on the visible CCD and the infrared spectrometer, both described by Rayman *et al.*¹ A conflict arises with the acquisition of science data with MICAS, as it is used to provide celestial inertial reference as well. As a result, MICAS will not be available for science data acquisition at all times throughout the encounter.

Because a comet encounter was not part of the primary mission, and funding and time were highly constrained during development, DS1 does not include dedicated shielding. Modeling of the dust environment based on fits to photometric data suggests that with an impact parameter of 2000 km, $\sim 10^2$ impacts of particles exceeding 40 μm in radius may be expected. More than half of those impacts are with particles under 200 μm in radius. This significant hazard cannot be reduced substantially without guaranteeing that the science objectives will not be achieved, so the risk is accepted.

The difficulty in developing an accurate radiometric navigation solution with the IPS operating, including the effects of the 1% - 2% uncertainty in the acceleration from the IPS, could cause targeting errors that are unacceptable, so impulse power thrusting to conserve

hydrazine will be terminated 7 days prior to the encounter; control of all three axes then will be accomplished with the RCS. The IPS still will be operated for acquisition of Borrelly images for optical navigation. The spacecraft angular deadbanding rates are lower in TVC mode, which uses a proportional controller instead of an impulsive “bang-bang” controller as with the RCS, so the IPS will be on whenever distant observations are attempted.

Dedicated trajectory correction maneuvers (TCMs) will be executed during the final two weeks before closest approach. The IPS will be used when there is time, but the achievable accelerations are low enough that terminal TCMs may require the RCS. Indeed, the dominant term in the hydrazine budget is the allowance for a TCM total of 10 m/s, consuming 2.0 kg.

Because there are spacecraft attitudes that are not safe, a strategy that biases the encounter targeting was adopted. The removal of the bias will require TCMs near the center of the space of allowed attitudes. In this way, each TCM will have a high probability of not requiring decomposition into segments whose vector sum accomplishes the required maneuver at the expense of greater hydrazine expenditure.

Science data acquisition at Borrelly will begin with PEPE switching to its high-rate encounter mode at 12 hours before closest approach (CA) to the nucleus.

The location of the nucleus will have to be inferred from the observations (both from Earth and from the spacecraft) of the coma. Even at CA - 1 day, the nucleus will be a small fraction of one MICAS 13- μ rad visible CCD pixel, and locating it in the presence of the confusion from the coma will be difficult. This complicates both the trajectory targeting and MICAS pointing strategies.

The relative photometric properties of the nucleus and the coma are highly uncertain, so when the nucleus will be detectable is very difficult to predict. With an approach phase angle of 91°, the effects of the nucleus’ shape and self shadowing and the presence and nature of jets significantly compound this problem. Acquiring a 50-pixel image will require imaging at approximately CA - 7 minutes (m). The *a priori* ephemeris will not be adequate to achieve

this, and this will not allow time for mosaicking, so autotracking will be required.

The spacecraft does not have the capability to adjust autonomously the integration time for the visible or infrared measurements. To finalize the selection of integration times, the spacecraft will acquire images at about CA - 11 hours and turn back to point the HGA to Earth to return them. The attitude in which MICAS is pointed at Borrelly permits only very low-rate communications with Earth through an LGA. The structure of the encounter sequences is such that if the rapid analysis of these images suggests that the baseline integration times are not suitable, one of two alternative sets, that will already be on board, could be selected with only one command. Within any one set, a range of integration times will be used. It is clear that with a nuclear rotation period of 25 hours, the selection of integration times based on data from CA - 11 hours, even if those data were unambiguous, would not ensure accurate imaging.

The time of closest approach was selected to allow this return of pre-encounter images and subsequent final commanding to occur during an overlap in coverage between the 70-m stations at the Goldstone and Canberra Deep Space Communications Complexes. Based on that criterion, encounter targeting controlled closest approach to be at approximately 22:30 UTC on the spacecraft.

Following the final high-rate communications, the spacecraft will turn to point MICAS to a star near Borrelly. Locking to this star will allow ACS to remove any errors accumulated during the turn and to continue reducing the errors in the IMU biases estimates. At about CA - 80 m, the spacecraft will begin its first MICAS science observations, with a set of infrared spectra. They will be followed by a turn to another star for a final attitude reference lock. The location of this star was chosen so that pointing to it will achieve a spacecraft attitude comparable to one part of the way through the final encounter sequence, thus reducing the effect of inertially fixed attitude errors that arise from the uncertainty in roll around the MICAS boresight.

To give ACS the best chance to estimate IMU biases, the spacecraft will be locked on this final reference star until about CA - 35 m. It

will then turn to place the predicted position of the nucleus in the MICAS visible CCD field of view and begin an imaging campaign. Each image will be delivered to software that is part of the autonomous optical navigation system,¹¹ some of which was uplinked in March 2001, with the rest having been used for other functions earlier in the mission. The software will search for an object that satisfies certain criteria associated with the nucleus, including maximum and minimum size (in two linear dimensions, not in angular extent), brightness, and positional consistency among images.

The difficulty of identifying the nucleus is exacerbated by significant scattered light in MICAS images. The scattered light problem was studied in detail during the primary mission; it is well understood and easily correctable in future versions of an instrument of a similar design. It made the recovery from the loss of the SRU more difficult, and it will make ground-based analysis of pre-encounter images for navigation and selection of final integration times more difficult. The scattered light significantly complicated the design of the software for the autotracking, and the extra constraints that had to be included to prevent tracking stray light raise the risk that the nucleus itself will not be tracked.

The coordinates of the object identified as the nucleus will be delivered to the core of the autotracking system known as RSEN (reduced state encounter navigation). RSEN will solve for the position of the nucleus relative to the spacecraft; in this reduced state, heliocentric positions are irrelevant. It also will solve for a drift of each IMU bias away from ACS' final estimate, which will be frozen when it stops tracking the reference star.

While ACS points to the predicted location of the nucleus, compensating for IMU biases with its last estimated values, RSEN can correct for biases that change linearly with time. RSEN's task will be complicated by the uncertainty of the cometary ephemeris. RSEN will weight early data to its estimation of IMU biases and later data to its estimation of the position of the nucleus. With a targeted impact parameter of 2000 km, along-track errors will not be significant at the beginning of this imaging session. In addition to the importance of selecting good integration times for science

imaging, integration times are important for RSEN; if the nucleus is not located and used in RSEN's filter, the solutions for IMU biases and nuclear ephemeris will suffer.

Transferring an image from the focal plane to the spacecraft computer consumes about 20 s. With another 10 s allocated for the subsequent processing of the image, images will be acquired at intervals of 30 s.

During the first 20 minutes of imaging, data will be fed to RSEN, but its solutions will not be used. This will give its filter time to converge without affecting spacecraft attitude until autotracking is required.

Two very short mosaics are planned, each containing only two attitudes away from the nominal. The mosaics will reduce the data available to RSEN in the nominal case, so their duration will be kept to a minimum. Beginning at about CA - 22 m and CA - 13 m, each mosaic will last about 3 minutes and may help in reconstructing the sequence of events should the nucleus fail to be captured in the other images.

Because of the uncertainty in the ephemeris, the time of closest approach will remain uncertain long past the final pre-encounter high-rate telecommunications session. Therefore, the sequence that contains all the commands near closest approach will be activated by AutoNav based on its estimate of the time of closest approach. At CA - 9 m AutoNav will be commanded to compute the time to closest approach and activate this final sequence when it is CA - 7 m. A back-up activation will be included in the event that AutoNav has not observed the nucleus and thus does not have a good solution.

Monte Carlo simulations confirm that the probability of the nucleus leaving the field of view rises dramatically during the final minutes. Imaging will be attempted until CA - 112 s, by which time the nucleus is likely to be out of the MICAS FOV. At that point, the attitude priority is shifted to the PEPE measurements in order to keep the smallest pixels in PEPE's 2.8π sr FOV aligned with the ram direction. Stopping the tracking for MICAS pointing is also important to allow the spacecraft to achieve a constant angular rate. This will reduce RCS activity near closest approach to minimize the interference of

hydrazine decomposition products with PEPE measurements and the interference of RCS solenoid activity with IDS magnetic field measurements.

PEPE and IDS will continue acquiring data past closest approach and through the subsequent turn to the reference star that allows the HGA to be Earth-pointed. No outbound remote sensing data will be collected for a variety of reasons including the long time it would take to rotate the spacecraft to repoint MICAS, the hydrazine cost of subsequent maneuvers, the continuously degrading pointing accuracy in relying on IMUs for so long, and the preference to fill the limited data storage volume with the higher probability of success inbound data.

IN-FLIGHT ENCOUNTER TESTS

Testing of sequences for the comet encounter depended upon extensive use of the DS1 testbeds at JPL. The testbeds could not reproduce all flight-system behavior, particularly spacecraft dynamics, so several tests with the spacecraft were conducted.

Limited project funding effectively prevented a modification of the software to allow ACS to track the comet nucleus in the same way it tracks a reference star. Therefore during encounter, the spacecraft will have to rely on the IMUs for pointing. These devices however drift with a character that is difficult to model accurately for individual events. Indeed, there is a significant random-walk component to the drift. In addition, the substantial filtering of the raw IMU output on board and limited telemetry complicates the inferences of the underlying behavior.

On May 1, the spacecraft used RSEN to track Jupiter from a range of 5.5 AU. Although the target's motion was negligible during the 2.5-hour test, it provided an opportunity to exercise the software on the spacecraft with actual IMU biases and actual MICAS images. Jupiter was about 13 pixels across, approximately the size of the nucleus near the beginning of the imaging following the last attitude update from a reference star.

Simulations of the encounter were conducted on the spacecraft on May 8 and June

28. These activities afforded an opportunity to verify processor loading, timing, spacecraft dynamical performance, and more with the sequences that were being planned for use at encounter. The sequences were not finalized by the time of these in-flight tests, but the simulations provided valuable reference data that contributed to the ongoing development.

The ephemeris of a fictitious comet (comet Spoof) was loaded, and the spacecraft behaved as if it were encountering that body. The plan was to intercept each image file before it was delivered to the image processing software and "paint" a nucleus and cosmic rays on it. The file was then to be returned to the data processing stream which would be presented with a target that changed in size and illumination phase in a realistic way. The same system was used regularly for testbed testing.

This scheme allowed much of the encounter to be tested, but each in-flight encounter rehearsal was compromised by one error in the simulation system. In the first one, a mistake in a parameter prevented the nucleus from being added to each image from the camera. Instead, the synthesized nucleus image replaced the camera's image. This prevented the rest of the software from seeing images with realistic scattered light. In the second test, a different error prevented the synthesizer from delivering the images back to the data processing stream. As a result, RSEN never received any data to process. It correctly declined to update the attitude, and the back-up activation for the final sequence containing commands for the time near closest approach received an unplanned test.

Most of the objectives of the in-flight rehearsals were achieved. It was decided not to conduct another simulation in order to capture the remaining objectives. The work required to prepare a rehearsal on the spacecraft was significant for the small team. In addition, the hydrazine expenditure and the additional spacecraft risk inherent in any such activity were deemed to be too high for the few remaining objectives. As a result, subsequent testing was confined to testbeds.

CONCLUSION

The encounter with comet Borrelly will be the final major activity of a mission that has

overcome numerous significant challenges. Reaching the comet and preparing for the encounter are not simply a continuation of the activities of the successful primary mission but rather are work of an entirely different ilk. The encounter would present important challenges even for a spacecraft built for the purpose, and the problems faced by DS1 are still greater. The benefits to science and to future comet missions in development and in flight from DS1's attempts to investigate comet Borrelly make the effort worthwhile given the very low cost of the extended mission. Because the project's resources have been quite limited, careful decisions in the management of risk have been an important ingredient in the success of difficult operations to date. Still, NASA's and JPL's acceptance of significant risk for the encounter has been an essential part of the planning of the conclusion of the Deep Space 1 Extended Mission.

ACKNOWLEDGMENTS

The members of the DS1 mission operations team are gratefully acknowledged for their dedication and excellent work, upon which this overview is based.

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

1. Rayman, Marc D., Philip Varghese, David H. Lehman, Leslie L. Livesay, "Results from the Deep Space 1 Technology Validation Mission," *Acta Astronautica* **47**, p. 475 (2000).
2. Deep Space 1 Technology Validation Reports, JPL Publication 00-10 (Jet Propulsion Laboratory, Pasadena, CA, October 2000).
3. Rayman, Marc D. and Philip Varghese, "The Deep Space 1 Extended Mission," *Acta Astronautica* **48**, p. 693 (2001).
4. Morabito, D., S. Shambayati, S. Finley, D. Fort, J. Taylor and K. Moyd, "Ka-band and X-band Observations of the Solar Corona Acquired During Superior Conjunctions Using Interplanetary Spacecraft," Proceedings of the Seventh Ka-band Utilization Conference, Genoa, Italy, September 26-28, 2001.
5. Rayman, Marc D. and Steven N. Williams, "Design of the First Interplanetary Solar Electric Propulsion Mission," submitted for publication.
6. A'Hearn, Michael F., Robert L. Millis, David G. Schleicher, David J. Osip, and Peter V. Birch, "The Ensemble Properties of Comets: Results from Narrowband Photometry of 85 Comets, 1976-1992," *Icarus* **118**, p. 223 (1995).
7. Lamy, Philippe L., Imre Tot, and Harold A. Weaver, "Hubble Space Telescope observations of the nucleus and inner coma of comet 19P/1904 Y2 (Borrelly)," *Astronomy and Astrophysics* **337**, p. 945 (1998).
8. Young, David T., Jane E. Nordholt, John J. Hanley, "Plasma Experiment for Planetary Exploration (PEPE)" in Deep Space 1 Technology Validation Reports, JPL Publication 00-10 (Jet Propulsion Laboratory, Pasadena, CA, October 2000).
9. Brinza, D. E., M. D. Henry, A. T. Mactutis, K. P. McCarty, J. D. Rademacher, T. R. van Zandt, P. Narvaez, J. J. Wang, B. T. Tsurutani, I. Katz, V. A. Davis, S. Moses, G. Musmann, F. Kuhnke, I. Richter, C. Othmer, and K.-H. Glassmeier, "An overview of results from the ion diagnostics sensors flown on DS1," AIAA paper 2001-0966, 39th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, January 8-11, 2001.
10. Richter, I., D. Brinza, M. Cassel, K.-H. Glassmeier, G. Musmann, C. Othmer, K. Schwingenschuh, B. Tsurutani, "First Direct Magnetic Field Measurements of an Asteroidal Magnetic Field: DS1 at Braille," *Geophysical Research Letters* **28**, p. 1913 (2001).
11. Desai, S., D. Han, S. Bhaskaran, B. Kennedy, T. McElrath, G. W. Null, J. E. Riedel, M. Ryne, S. P. Synnott, T. C. Wang, R. A. Werner, E. B. Zamani, "Autonomous Optical Navigation (AutoNav) Technology Validation Report" in Deep Space 1 Technology Validation Reports, JPL Publication 00-10 (Jet Propulsion Laboratory, Pasadena, CA, October 2000).